Magnetometry with Ensembles of Nitrogen Vacancy Centers in Bulk Diamond

by

Carson Teale

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2015

Distribution A: Public Release.

Author . . . .
Department of Electrical Engineering and Computer Science
August 28, 2015

Certified by. . . . .
Dirk Englund
Assistant Professor
Thesis Supervisor

Certified by. . . . .
Danielle Braje
Research Scientist
Thesis Supervisor

Accepted by . . . .
Professor Leslie A. Kolodziejski
Chairman, Department Committee on Graduate Theses

Magnetometry with Ensembles of Nitrogen Vacancy Centers in Bulk Diamond

by

Carson Teale

Submitted to the Department of Electrical Engineering and Computer Science
on August 28, 2015, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

This thesis summarizes experiments conducted to develop a high sensitivity vector magnetometer using nitrogen vacancy (NV) centers in a bulk diamond sample. This project began by analyzing the sensitivity of a single NV orientation using a continuous wave electron spin resonance approach. A protocol for determining the diamond's orientation was developed to map vector magnetic field readings in the diamond reference frame to the lab frame. Preliminary vector field measurements and differential vector measurements were performed. Although these showed promising results, significant instrument and ambient magnetic noise limited the achievable sensitivity. A new frequency locking measurement technique was developed to allow for simultaneous measurements between two separate sensors for future differential experiments. This technique provides a host of other benefits including much improved dynamic range and steady-state immunity to fluctuations in linewidth and contrast.

Thesis Supervisor: Dirk Englund
Title: Assistant Professor

Thesis Supervisor: Danielle Braje
Title: Research Scientist

Acknowledgments

I'd like to thank both of my advisors, Danielle Braje and Dirk Englund, for giving me the opportunity to do this research and providing me with guidance along the way. I'd also like to thank Hannah Clevenson who worked with me on many of these experiments, Michael Walsh for all the help he provided, and all of the other members of the Quantum Photonics Group.

Contents

6 Conclusion 53

List of Figures

11

Chapter 1

Introduction

The precision measurement of magnetic fields has many practical applications ranging from navigation to object detection to medical imaging [19, 8, 18]. However better sensitivities and smaller device sizes will be required for many future applications. Most low frequency, medium sensitivity applications use fluxgate magnetometers while applications requiring higher sensitivities would use SQUID or atomic magnetometers. Fluxgate magnetometers are low-cost, vector sensors with sensitivities as low as 10 pT for low frequency fields; however thermal effects lead to drifts of around 0.1 nT/^{o}C [21]. Atomic magnetometers have achieved sensitivities below 1 fT/$\sqrt{\text{Hz}}$ with a package size of less than 0.5 cm^3 but are somewhat difficult to integrate with standard CMOS processes because of the challenges involved with fabricating the vapor cell [16]. Nitrogen vacancy centers are a color center in diamond that behave similarly to an isolated atom but are confined in a solid-state lattice. An NV based magnetometer, like an atomic magnetometer, has the potential to combine high sensitivity and low drift but in a much smaller package size. An NV magnetometer is also better suited for vector measurements than most other technologies because of the four well-defined NV orientations in the diamond lattice

1.1 NV Basics

The nitrogen vacancy center is a point defect in diamond where a nitrogen atom and neighboring vacant lattice site replace two carbon atoms in the diamond crystal. The NV can exist along any one of the four crystallographic axes of the diamond lattice (see Fig. 1-1a). The NV exists in two fluorescent charge states, negatively charged (NV^-) and neutral (NV^0) [12]. All of the experiments described here are performed with the negatively charged states. The energy level diagram for NV^- is shown in Fig. 1-1b). The negatively charged center has six electrons, of which two are unpaired and form a triplet ground state and optically excited state as well as several dark singlet states which relax back down to the triplet ground state [5]. The zero phonon line (ZPL) between the triplet ground and excited states is at 637 nm; however, excitation with higher energies is possible through phonon mediated transitions [23]. The most remarkable feature of the NV is its spin dependent fluorescence and optical polarizability at room temperature. When the NV is optically excited from the ground state to the excited state, it decays back down to the ground state via two pathways. It can either cycle back to the ground state through the emission of a 637 nm photon in which case the spin of the system is conserved, or it can go through a non-radiative, non-spin conserving decay path to the singlet state manifold. The transition rate from the $m_s = \pm 1$ excited state to the dark state is faster than that of the $m_s = 0$ excited state. This results in the spin non-conserving transition occurring more often when the system is in the $m_s = \pm 1$ state, meaning that the spin more often stays in the $m_s = 0$ state. Thus by continually exciting the NV for around 300 ns, the spin can be polarized to $m_s = 0$ with polarization around 80-90% [22]. This same effect allows for spin state fluorescent readout. Since the $m_s = \pm 1$ states go through the non-radiative transition more often than the $m_s = 0$ state, they emit less fluorescence when cycled through the excited state.

The structure of the NV's ground state energy levels can be modeled by the ground state Hamiltonian, which neglecting strain effects and hyperfine coupling to nuclear spin, is [6]:
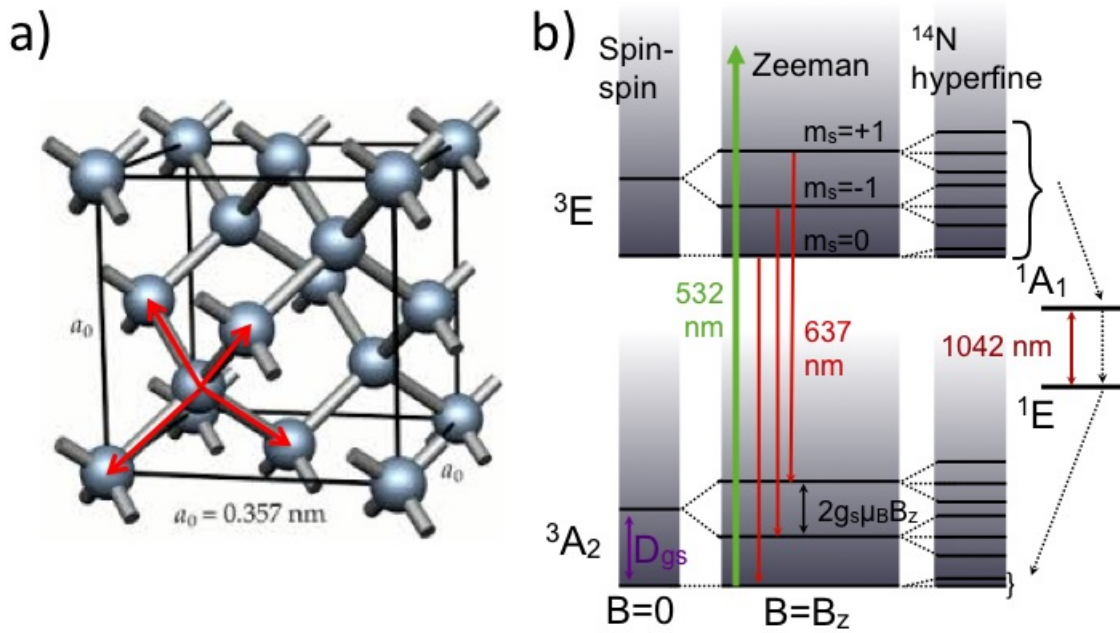
Figure 1-1: a) Diamond lattice structure. Red arrows show the crystallographic axes of the diamond, which are also the directions of the four possible NV orientations (the NV can be either aligned or anti-aligned with each of the four arrows). b) Energy level of the NV^-. The left most section shows the energy levels with no magnetic field present, the center section shows the splitting that occurs with a magnetic field, and the right most section shows the hyperfine structure of the energy levels.

$$H = hDS_z^2 + \gamma \mathbf{B} \cdot \mathbf{S} \qquad (1.1)$$

D is a zero field splitting (ZFS) parameter of around 2.87 GHz, $\mathbf{S}$ is the vector spin operator consisting of S_x, S_y and S_z Pauli matrices, and $\gamma = g_e \mu_b / h \approx 28$ Hz/nT is the gyromagnetic ratio of an electron in the NV. The ZFS between the $m_s = 0$ and $m_s = \pm 1$ energy levels comes about because of spin-spin interactions where electron spins pointing in the same direction have higher energy than when the two spins are in opposite direction. The energy difference depends on the extent to which the two electron wavefunctions overlap which is itself a function of the lattice structure where the NV resides. Temperature shifts can deform the diamond lattice resulting in shifts in the zero field splitting [1]. The $m_s = \pm 1$ states have equal energies with no applied magnetic field. For small magnetic fields, these two states split apart at a rate of $\Delta \nu = 2\gamma B_z$ where B_z is the projection of the magnetic field along the NV axis. At room temperature the excited state has a similar fine structure with a zero field splitting of 1.41 GHz.

Fig. 1-1b) shows the hyperfine structure of the NV ground state. Most of the NVs in this sample are composed of ^{14}N with spin 1. When hyperfine coupling between the electron and nuclear spin are taken into account, the $m_s = \pm 1$ states split into $|m_s = \pm 1, m_I = -1\rangle$, $|\pm 1, 0\rangle$ and $|\pm 1, +1\rangle$. With $m_s = 0$ there is no hyperfine coupling; however it still splits into two states: $|m_s = 0, m_I = 0\rangle$ and $|0, \pm 1\rangle$ which has higher energy due to the alignment of the two nuclear spins [11].

1.2   NV as a Vector Magnetometer

By measuring the Zeeman shifts in the energy levels of the $m_s = \pm 1$ states, the projection of a magnetic field onto the NV axis can be found. The simplest way to measure these energy levels is by observing electron spin resonance (ESR) where the NV is continuously illuminated with laser light (typically 532 nm) and microwaves [7]. The microwave frequency is swept in a range near the zero field splitting (2.87 GHz) while fluorescence from the diamond is collected. When the microwaves are not on
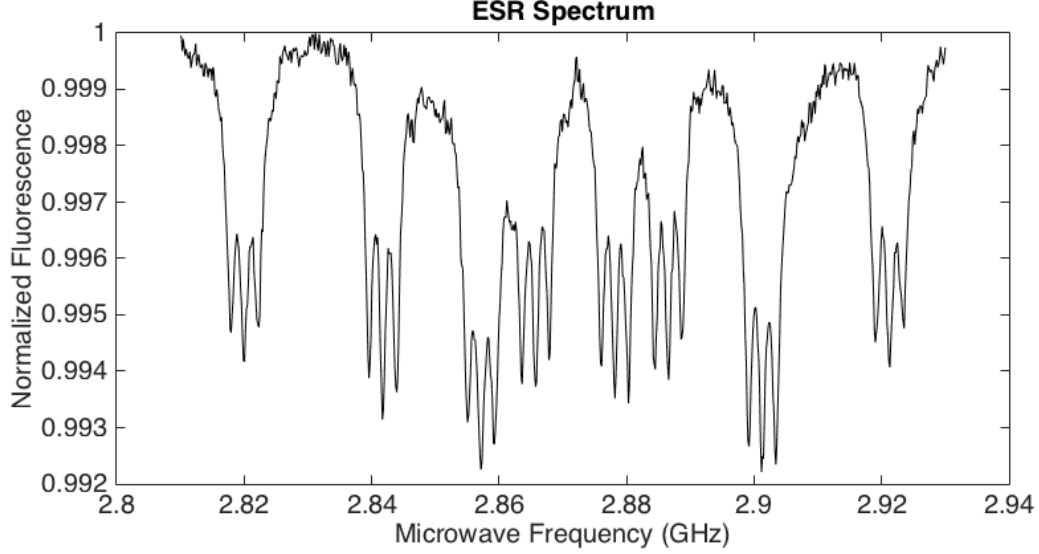
ESR Spectrum

Figure 1-2: Electron spin resonance (ESR) spectrum with a magnetic field oriented such that the projection onto each axis is different, resulting in eight separate resonances ($m_s = 0 \rightarrow \pm 1$ for each of the four axes).

resonance with the $m_s = 0 \rightarrow \pm 1$ transition of the NV, the green laser continually polarizes the NV, keeping it in the $m_s = 0$ state. When the microwaves are on resonance with one of these transitions, the NV will occasionally be excited to the $m_s = \pm 1$ state at a rate determined by the relative rabi frequencies of the microwave and optical transitions. Now, because the NV is sometimes in the $m_s = 1$ or $m_s = -1$ state and not always the $m_s = 0$ state, the fluorescence is diminished compared to the off resonant case. Thus, as the microwave frequency is swept, dips in the fluorescence appear at the NV resonances. In a bulk, single crystalline diamond, the NVs typically exist in all four possible orientations. With zero applied magnetic field, all energy levels are degenerate and only one resonance line appears at the ZFS. An applied field breaks this degeneracy and up to eight separate lines may be observed, as in Fig. 1-2. In this ESR spectrum, each large resonance line is composed of three smaller lines which correspond to the three energy levels from hyperfine coupling to the nitrogen nuclear spin.

External magnetic fields can be measured by recording ESR spectra. First a static magnetic field is applied to the diamond to separate the NV resonances and then the frequency shift of each resonance is recorded when the unknown field is

17

applied. This frequency shift can be translated to a magnetic field projection and if the frequency shifts of all eight resonances (corresponding to the four NV orientations) are recorded, the vector magnetic field can be determined by knowing the unit vectors of the diamond lattice and the corresponding field projections.

1.3   CW Sensitivity

The shape of the resonance lines in the ESR spectrum can be approximated as Lorentzians where the intensity of the fluorescence is given by:

$$I(\nu) = R \left(1 - C \left(\frac{\Delta \nu^2}{(\nu - \nu_0)^2 + \Delta \nu^2} \right) \right) \tag{1.2}$$

Where I is the fluorescence intensity, R is the rate of photon collection, C is the contrast, ν is the applied microwave frequency, ν_0 is the resonant frequency, and $\Delta \nu$ is the linewidth of the transition. The number of photons collected follows a Poisson distribution with standard deviation scaling as the square root of the number of photons collected: $\delta I = \sqrt{I \Delta t}$ (assuming shot noise limited performance). If the measurement is shot noise limited, then the sensitivity η is given by [7]:

$$\eta = \delta B_{min} \sqrt{\Delta t} = \frac{h}{g_e \mu_B} \frac{\sqrt{I}}{\max \left| \frac{\partial I}{\partial \nu_0} \right|} \approx 4/3\sqrt{3} \frac{h}{g_e \mu_B} \frac{\Delta \nu}{C \sqrt{R}} \tag{1.3}$$

This derivation assumes that the measurement is biased at the microwave frequency ν that maximizes the slope of the intensity versus frequency. The sensitivity scales directly with linewidth and inversely with contrast and photon collection rate. To improve contrast, it is necessary to increase the power of the microwave fields to more quickly drive the transition from $m_s = 0 \to \pm 1$ and to increase photon collection rate, it is necessary to increase the laser pump power. However, both high laser power and high microwave power result in power broadening of the linewidth, so that optimal sensitivity is achieved through a compromise of increasing contrast and collection rate while avoiding power broadening.

1.4   Pulsed Measurements

External magnetic fields can also be measured using pulsed techniques, where neither the green pump laser nor the microwaves are applied continuously [20, 24, 10]. Instead, each is applied separately: a laser pulse first polarizes the NV to the $m_s = 0$ state, then a microwave pulse is used to induce a Rabi oscillation between $m_s = 0$ and $m_s = \pm 1$. For a Ramsey experiment, the duration of this pulse is chosen so as to drive the system to the superposition state $m_s = 1/\sqrt{2}(|0\rangle + |1\rangle)$ (a pi/2 pulse). A relative phase proportional to the energy difference between $|0\rangle$ and $|1\rangle$ (which is a function of magnetic field) is accumulated during a free precession interval, such that the state becomes: $1\sqrt{2}(|0\rangle + |1\rangle e^{i(E_1 - E_0)\tau/\hbar})$. Another pi/2 pulse about the same axis as the first pi/2 pulse rotates the state off of the equator of the Bloch sphere. Finally the state is projected back on to the $|0\rangle$ and $|1\rangle$ basis when a second laser pulse is used to read out the spin state by measuring the spin dependent fluorescence. This type of pulsed measurement typically results in better sensitivity than the continuous wave approach because the contrast in fluorescence between the two states can be maximized and linewidth broadening due to high microwave and/or pump laser power can be eliminated. AC magnetometry can be performed by using more complex pulse sequences that filter out magnetic noise at frequencies far from the frequency of interest.

1.5   Advantages of the NV as a Magnetometer

The splitting of the NV's $m_s = \pm 1$ energy levels as a function of magnetic field can be calculated directly from the Hamiltonian, which is written in terms of physical constants. This contrasts with fluxgate magnetometers which must be calibrated and whose response to a magnetic field varies from device to device [21]. This also means that as a fluxgate ages, it may need to be re-calibrated as the device deteriorates with time. The ground state energy levels of every NV center respond in the exact same way to an applied magnetic field and this response does not change over time. This is

similar to atomic magnetometers whose response is also tied directly to fundamental constants and whose response is identical from atom to atom or device to device. However these atomic devices require a relatively large gas chamber and generally cannot be fabricated using standard lithographic techniques. They also inherently measure a scalar field, resulting in additional challenges for vector measurements. NV centers are already contained in the diamond lattice which can be grown and/or cut to a very small size [9]. The diamond can then be attached to a chip which contains the other necessary components of the magnetometer fabricated through standard silicon processing.

For many applications, it is desirable to record the vector magnetic field rather than just the magnitude. Fluxgate magnetometers achieve vector sensing by fixing three scalar magnetometers orthogonal to one another [3]. The angular accuracy of this approach depends on the precision of the manufacturing process. Another advantage of an NV based magnetometer is its inherent vector nature. NVs in a single crystalline diamond exist along all four orientations of the tetrahedral diamond lattice These four axes provide an over-complete basis where the direction of each unit vector is precisely known and determined purely by geometry.

The rest of this document describes the development of experiments and techniques for the purpose of developing an NV based, vector magnetometer that uses these desirable features. A great deal of research focused on these same goals had been performed prior to the work described here [4, 17, 15]. The main challenges addressed here are mapping the diamond lattice frame to the lab coordinate system, the design of a uniform field, permanent magnet array, and a scheme for continuous wave, simultaneous, vector measurement.

Chapter 2

Measurement of Single Axis Sensitivity

Before attempting a vector field measurement, the sensitivity of a measurement along a single axis was optimized.

2.1   Experiment Setup

Fig. 2-1 shows a diagram of the basic experimental setup. A type IIa, commercial, optical grade, CVD grown diamond with a $\langle 100 \rangle$ cut and dimensions of $3 \times 3 \times 0.3$ mm^3 was used for all of the following experiments. The diamond was electron irradiated and subsequently annealed with a dosage resulting in 0.1 ppm NV centers. This concentration of NVs represents a good compromise between maximizing signal strength while preventing coupling between adjacent NVs which would degrade overall spin coherence [20]. A 532 nm laser is spatially filtered and used to illuminate the diamond sample. The sample is mounted on a copper pedestal which is secured to a water cooled platform for maintaining temperature stability. The stage is positioned inside a three axis Helmholtz coil array such that the diamond sample is centered in the middle. The coils are powered using low noise current sources capable of delivering up to 1 A, which corresponds to a magnetic field of around 30 G. Two aspheric condenser lenses are used to direct the fluorescence from the diamond to a
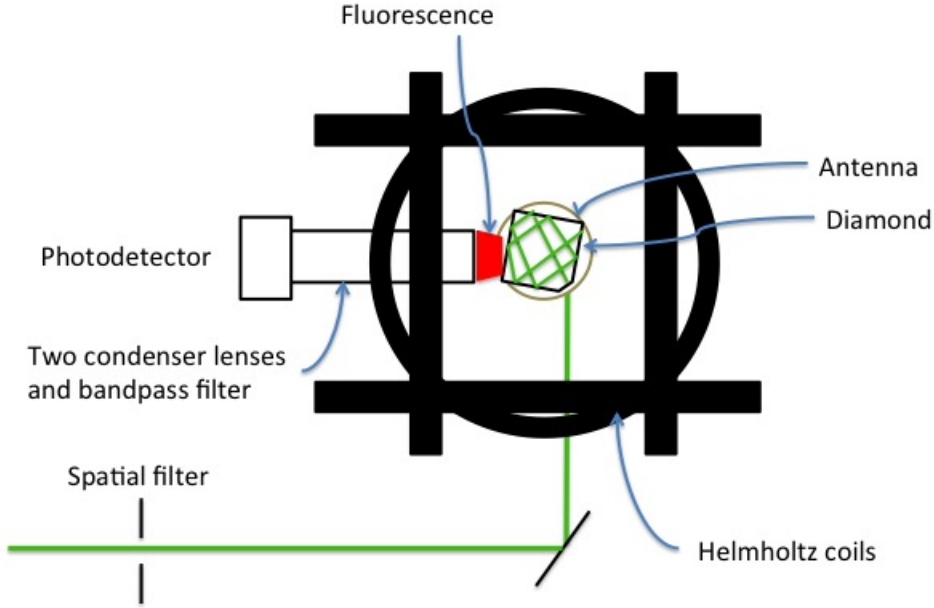
Figure 2-1: Top view diagram of basic experimental setup

large area photodetector. A 662-800 nm bandpass filter is used to cut out any leftover 532 nm pump light as well as NV^0 fluorescence (this filter also blocks infrared light from the laser at 1064 nm and 1042 nm light from the radiative, intersystem crossing transition). A 3 mm diameter loop of wire suspended 5 mm above the diamond functions as a microwave antenna. It is connected to an RF signal generator and driven at around 40 dBm of microwave power.

For a shot noise limited measurement, the signal to noise ratio goes as the square root of the number of photons which is proportional to the number of emitters (NVs). Thus it is beneficial to address as many of the NVs in the diamond sample as possible. Prior to this work, most experiments using bulk diamond samples used a single pass geometry where the green excitation laser passes once through the diamond. For this experiment and all others described in this document, a light trapping diamond waveguide (LTDW) structure was used to increase the path length of the green laser
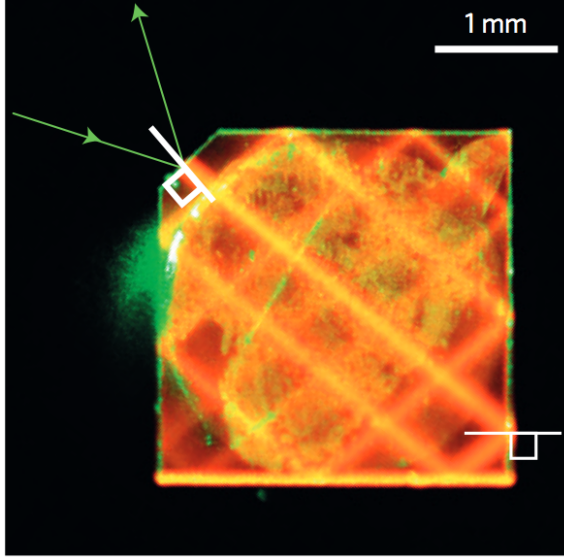
Figure 2-2: Photograph of light-trapping-diamond-waveguide (LTDW) structure. Green laser light enters at the input facet and is trapped through total internal reflection (TIR). The pattern is orange and not green because while the green light is trapped by TIR, the orange fluorescence of the NV is emitted uniformly in all directions, allowing most of it to escape.

within the diamond in order to excite more NVs. This a very simple, yet elegant structure consisting of the diamond sample described above where each of the six surfaces is polished and a 500 μm facet is cut at an angle of 45^o at one of the corners. The laser light enters through this facet and is then confined inside the diamond by total internal reflection until it either hits another corner or comes back out the facet. Depending on the entrance angle and position of the laser light, this structure can increase the path length by several orders of magnitude as compared with a single pass geometry.

2.2   Experiment Description

The experiment is very similar to the description of continuous wave electron spin resonance (CW ESR) experiments described in chapter 1. One addition is the use of a lock-in amplifier to increase the signal to noise by eliminating detector noise. The lock-in amplifier provides a reference oscillator which was set to a frequency around

1 kHz and connected to the frequency modulation input of the frequency generator. The modulation depth was set at 1.2 MHz. This was chosen to be close to the width of the resonance to generate the largest lock-in signal. The frequency modulation of the microwave signal is converted to an amplitude modulation of the fluorescence due to the frequency dependent shape of the ESR spectrum. The fluorescence signal goes to the input of the lock-in amplifier where the signal is demodulated. In general, a lock-in amplifier performs a homodyne type measurement where the input signal is mixed with the reference oscillator to produce sum and difference frequencies. The output of the mixer is then low-pass filtered with a selectable cutoff frequency, f_c, specified by a time constant, TC (for our lock-in $f_c \approx 1/(5TC)$). Frequency components of the signal near the reference frequency will be mixed down to DC and pass through the low-pass filter.

In this experiment, at a particular frequency point along the ESR curve, the amplitude modulation seen in the fluorescence scales linearly with the slope of the ESR curve at that point. The output of the lock-in is the envelope of the amplitude modulated fluorescence which is proportional to the derivative of the ESR curve. Any frequency components of the photodetector signal which are not close to the reference frequency, are filtered out. This mitigates the effects of any laser intensity noise and stray collected light not near the reference frequency.

Fig. 2-3a) shows a typical ESR curve along with the corresponding lock-in output b). In this case, a static magnetic field was aligned with one of the NV orientations, resulting in an equal projection onto the other three orientations, as shown in Fig. 2-3c). The goal of the measurement is to determine the magnetic field projection along one of the NV orientations by measuring the change in the frequency splitting between the two transitions, $m_s = 0 \rightarrow \pm 1$. In order to measure frequency shifts, the lock-in signal (Fig. 2-3b) is recorded and the slope of the signal at the center of each of the two resonances of interest is determined. Once the slopes are known, the frequency of the microwave signal generator (ω_{ODMR}) is fixed at the center of one of the resonances and the lock-in signal is again recorded. If there is no change in the magnetic field, the lock-in output will remain at zero. However, if the magnetic

24

field projection does change, the frequency of the resonance will shift while the signal generator's frequency stays put. Now a different part of the ESR curve with a different slope is being sampled. The change in the ESR slope results in a change in the lock-in output. This shift in the lock-in output voltage can be converted back to a frequency shift using the slope of the lock-in signal previously recorded. This procedure is shown schematically in Fig. 2-3d).

To measure the shift of both transitions ($m_s = 0 \rightarrow \pm 1$), the microwave source hops back and forth between the two original resonant frequencies (ω_+ and ω_-), recording the lock-in value at each. The frequencies of the transitions can be written in terms of the magnetic field projection and the temperature: $\omega_\pm \approx D + \beta_T \Delta T \pm \gamma B_z$. Where $D \approx 2.87$ GHz is the zero field splitting, $\beta_T \approx -74$ kHz/K (at room temperature) is an empirical temperature dependence factor, ΔT is the change in temperature, γ the gyromagnetic ratio, and B_z the projection of the magnetic field [1]. By subtracting these two equations, we can independently solve for the magnetic field projection: $B_z \approx (\omega_+ - \omega_-)/(2\gamma)$. By adding the equations we can solve for the change in temperature: $\Delta T \approx ((\omega_+ + \omega_-) - 2D)/\beta_T$.

2.3   Results

To estimate the sensitivity of the instrument, changes in the ω_+ resonance frequency were recorded over a 100 second duration with no applied field. The shifts in the resonance frequency during this time are either caused by ambient field noise or are not true frequency shifts, but instead laser or microwave power fluctuations that result in changes in the lock-in voltage output. In either case, this represents a typical noise floor and sets the lower limit on the size of magnetic fields measurable with the instrument. Fig. 2-4a) shows the power spectral density of this noise trace. At a frequency of 1 Hz, the instrument has a noise floor of around 400 pT/$\sqrt{\text{Hz}}$, meaning that it can detect any signal with a power density larger than that value. Fig. 2-4b) shows the separation of magnetic field and temperature obtained by measuring the frequency shifts of ω_+ and ω_- while a square wave magnetic field is applied.
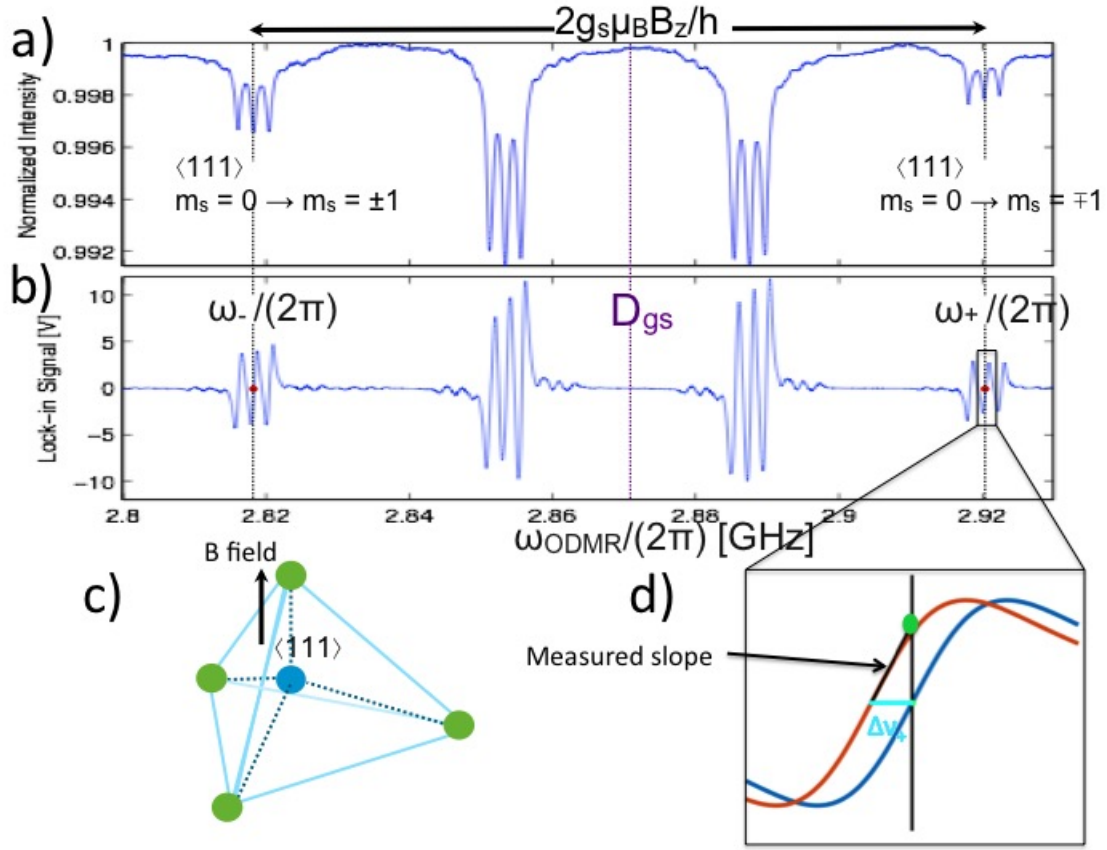
Figure 2-3: a) An ESR spectrum with a static magnetic field aligned with one of the four orientations, resulting in an equal projection along the other three. Here the change in the projection of the field along the orientation aligned with the static field is being measured (the frequency difference between the two outermost dips). Only four dips appear here, instead of eight, since resonances of three of the NV orientations are degenerate owing to an equal magnetic field projection. This also explains why the contrast of the middle two dips is about three times as large as the outer two. b) The lock-in output of the ESR spectrum in a) which is proportional to the derivative. c) A diagram of the four NV orientations where the blue dot represents a nitrogen atom and the green dots possible vacancy site locations (note that orientations with the nitrogen atom and vacancy center interchanged, also exist). d) Zoomed in on the lock-in output of a single resonance. The vertical black line is the fixed frequency of the microwave source. The blue curve is the original lock-in voltage vs. frequency. With a change in magnetic field, the curve shifts by $\Delta\nu_+$ to the red line and the lock-in output shifts from zero to the value at the green dot. The value of $\Delta\nu_+$ is recovered by using the previously measured slope value.
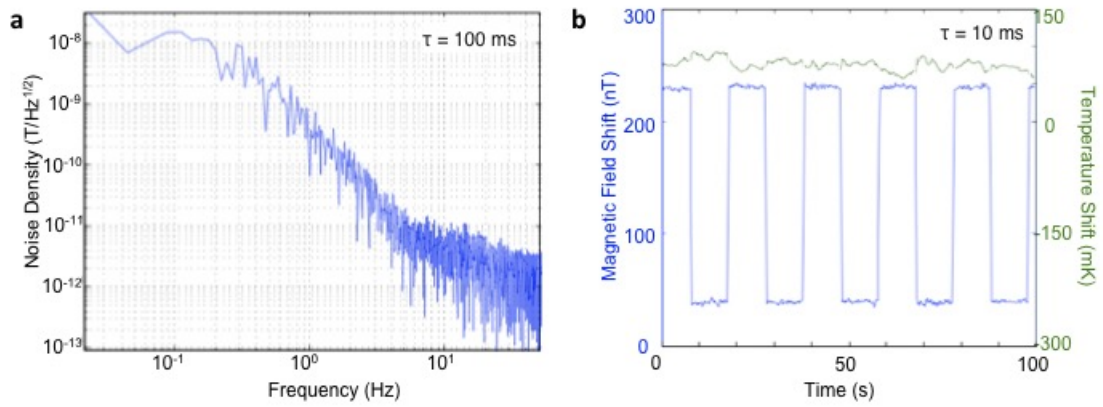
Figure 2-4: a) Power spectral density of noise trace taken over 100 seconds. The time constant of the lock-in amplifier was set to 100 ms. b) Magnetic field (blue) and temperature (green) with a 20 s period, 190 nT amplitude applied square wave. Here the time constant was set at 10 ms.

This level of sensitivity puts our device in the range of the most sensitive fluxgate and MEMs based magnetometers [8].

Chapter 3

Mapping of the Diamond Lattice to the Lab Frame

In order to apply magnetic fields at definite angles with respect to the diamond lattice, it was necessary to develop a procedure to find a rotation matrix between the diamond coordinates and the lab frame. The diamond is usually cleaved along its crystallographic planes, so by measuring the angles of the surfaces of the diamond it is usually possible to get a rough estimate of the diamond lattice orientation. However this is a very crude method and would not work for samples cut in a different manner. By applying known magnetic fields and recording the positions of the resonant frequencies in an ESR spectrum, it is possible to calculate a rotation matrix between the diamond and lab frames.

3.1 Symmetries of the NV Tetrahedral Orientations

As stated previously, NV centers can exist with their axis (defined as the line between the nitrogen atom and the vacant lattice site) along any of the four principle lattice vectors of the diamond crystal. In the bulk diamond sample used for these experiments, equal numbers of NVs exist along all of the possible orientations. The three-fold rotational symmetry plus the three-fold mirror symmetry about each of the four vectors of the tetrahedral structure means that in general there are 24 magnetic
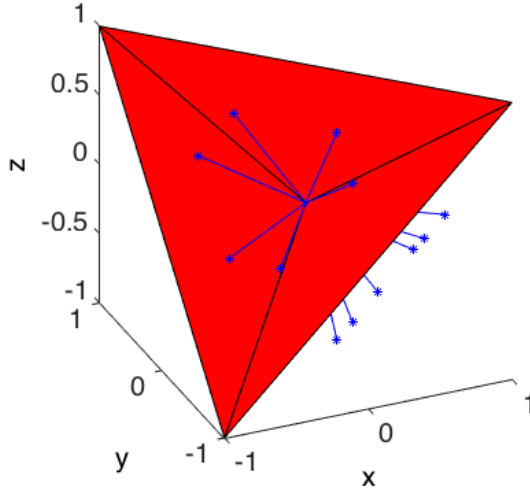
Figure 3-1: Tetrahedral structure shown in red with 24 possible vectors (blue lines) that result in the same projections (only 12 lines can be seen from this perspective, the other 12 are behind the red surfaces).

field vectors that produce the same ESR spectrum. Each of these 24 vectors have a different projection along each NV axis, but the ESR curves are the same since each of the four axes are identical and it is impossible to label which is which. This number is decreased when the magnetic field results in degenerate energies (equal projections) for multiple axes. For example, there are only four possible vectors that would result in the ESR curve shown in Fig. 2-3a) where the field is aligned to one of the diamond axes. Fig. 3-1 shows an example of 24 magnetic field vectors which all produce the same ESR spectrum.

The goal for a vector magnetometer would be to apply a static field to eliminate degeneracy between the four NV orientation so that each can be monitored independently, and then to determine the physical direction that each resonance corresponds to in the lab frame.

3.2 Procedure for Determining Lattice Orientation

One approach to breaking the symmetry is to apply at least three linearly independent, control magnetic field vectors (typically orthogonal vectors along the coordinate

axes defined by the Helmholtz coils: $\mathbf{B_x}, \mathbf{B_y}$, and $\mathbf{B_z}$) and use the resulting ESR spectra to determine the lattice orientation. For each field, the ESR spectrum is recorded and from this, a set of up to 24 vectors is calculated, where each vector produces the measured ESR spectrum. The algorithm then goes through these three lists of 24 vectors, looking at sets of three (one from each list) and finds sets where the three vectors are orthogonal and have the correct orientation with respect to one another. This is done by computing the dot product of the first with the cross product of the second and third $(\mathbf{B'_{x_n}} \cdot (\mathbf{B'_{y_m}} \times \mathbf{B'_{z_l}}))$ and comparing the result with the same computation performed on the known control fields in the lab frame (here primed vectors are in the diamond frame, unprimed are in the lab frame). This will result in twelve sets of vectors (due to the three-fold rotational symmetry about each of the four axes; reflection symmetries are eliminated by preserving the handedness of the control vectors). A rotation matrix between each of these twelve sets and the lab frame set of vectors is then calculated: $r_{ij_n} = \mathbf{B_i} \cdot \mathbf{B'_{j_n}}$ (i,j = x,y,z). Any of these twelve matrices can be used to transform back and forth between lab and diamond reference frames as long as consistency is maintained.

This same procedure can be used when an unknown static magnetic field ($\mathbf{B_s}$) of non-negligible strength is present. In this case, an ESR spectrum is first recorded with no control fields applied and a list of potential vectors is generated ($\bar{\mathbf{B}}'_s$). After this, the same procedure is performed with the three control fields as described previously. In order to find the sets of orthogonal vectors in the diamond frame, it necessary to first subtract the static field vector. So for each of the 24 possible vectors generated with the applied control field ($\mathbf{B'_{T,i_n}}$), the algorithm finds the static vector ($\mathbf{B'_{s_m}}$) such that $|\mathbf{B'_{T,i_n}}| = |\mathbf{B_i}| + |\mathbf{B'_{s_m}}|$ and subtracts it from $\mathbf{B'_{T,i_n}}$. This is done for each of the three sets and then the algorithm proceeds, as before by finding sets of these vectors which are orthogonal.

Once the rotation matrix is found in this manner, the same set of ESR spectra can be used to find the magnitude and direction of the unknown static field. The direction and magnitude of the control fields are now known in both frames, so to find the direction of the static field, the algorithm just searches through the $\bar{\mathbf{B}}'_s$ vectors
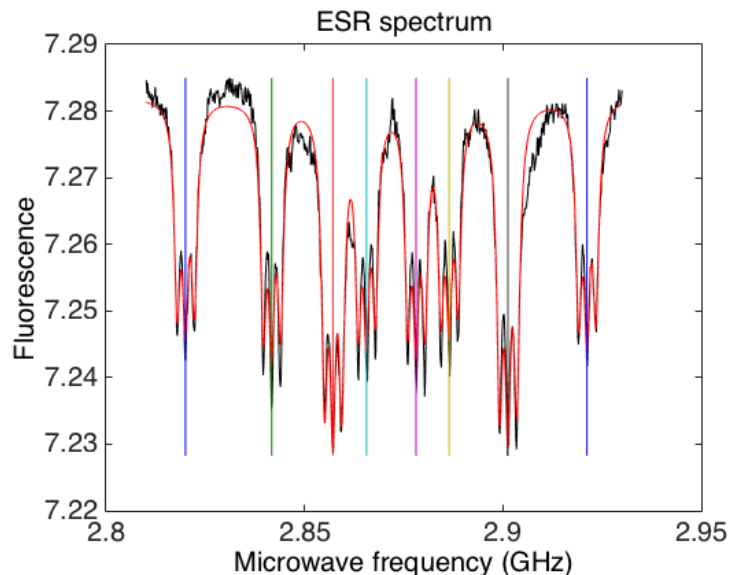
Figure 3-2: Measured ESR spectrum (black) and least-squares fit (red). Vertical lines show center frequencies of the eight resonances.

and $\bar{\mathbf{B}}'_{\mathbf{T}}$ vectors and finds the pair where the difference is equal to the control field ($\mathbf{B}'_{\mathbf{i}}$). This can be done for all three control fields to compensate for experimental imperfections.

3.3   Details of Algorithm Implementation

To extract the resonant frequencies from the ESR spectrum, the measured spectrum is fit to a function composed of eight sets of three Lorentzian functions ($(1 + ((f - f_0)/\Delta)^2)^{-1}$), representing the hyperfine structure of each electronic resonance. User input is used to seed Matlab's nonlinear least-squares fitting function which adjusts the contrast and frequency of each of the eight sets of Lorentzians to optimal values. Fig. 3-2 shows a typical measured spectrum and fit.

Once the resonant frequencies have been determined, they can be used to solve for a magnetic field. By inspecting Eq. 1.1, it is clear that the off-diagonal terms proportional to $B_\perp$, the transverse components of the magnetic field, can be ignored for small fields ($\gamma B_\perp \ll hD$). This is the same as assuming that the energy splitting of the $m_s = \pm 1$ states is proportional to the magnetic field projection (B_z). Following

this assumption, the system of equations to be solved is:

$$hD \pm \gamma \mathbf{B} \cdot \boldsymbol{\mu_i} = \pm \nu_i \tag{3.1}$$

Where $\boldsymbol{\mu_i}$ for $i = 1, 2, 3, 4$ are the four diamond unit vectors. Written in matrix form:

$$hD + \gamma \bar{\boldsymbol{\mu}} \mathbf{B} = \boldsymbol{\nu} \tag{3.2}$$

Where $\bar{\boldsymbol{\mu}}$ is an 8×3 matrix whose first four rows are the $\boldsymbol{\mu_i}$ vectors and whose last four rows are $-\boldsymbol{\mu_i}$ and $\boldsymbol{\nu}$ is a vector whose first four entries are the upper four resonant frequencies and whose last four entries are the lower four resonant frequencies. This system of equations can easily be inverted to solve for $\mathbf{B}$ given $\boldsymbol{\nu}$, though the answer will not be unique. Experimentally, the only challenge is that it is unknown which diamond axis ($\boldsymbol{\mu_i}$) corresponds to which frequency ($\boldsymbol{\nu_i}$). To account for this, the algorithm tries all possible combinations and picks the set with the smallest residuals from the least-squares inversion.

To get a more accurate value for the magnetic field vector, the result of the least-squares calculation is used to seed a nonlinear search which uses the full Hamiltonian (Eq. 1.1) to find a magnetic field vector that minimizes the difference between its resonant frequencies and the measured frequencies. This process uses Matlab's *fmincon* function. This function tends to get stuck at incorrect, local minima when not seeded with the results of the approximate linear least-squares solution and converges much more rapidly when seeded.

Once a possible magnetic field vector is found, the other 23 solutions are found by repeatedly reflecting the original vector about each of the planes defined by the lattice unit vectors, $\boldsymbol{\mu_i}$. This takes into account both the rotation and reflection symmetries of the tetrahedral structure.

Chapter 4

Preliminary Vector Field Measurements

Having analyzed the sensitivity of a single axis and developed a method to determine the diamond lattice orientation, initial vector field measurements were conducted. These experiments combined the techniques described in the previous two chapters to sequentially measure the magnetic field projection along each of the four NV orientations and then reconstruct the vector field, using the rotation matrix from the diamond frame to the lab frame.

4.1   Design of Uniform Field Permanent Magnet Array: Halbach Ring

To measure the resonant frequency shifts of each NV orientation separately, it is necessary to apply a relatively large (10 - 60 G) field in a direction that results in a different projection along each axis. Initially, this was done using the array of Helmohltz coils. However, the maximum field was then limited to around 30 G when the current sources were driven at their maximum value (around 1 A). For these higher current ranges, the sources have much larger current noise (up to 25 μA p-p which translates to around 80 nT between 0.1 and 10 Hz). To avoid this added
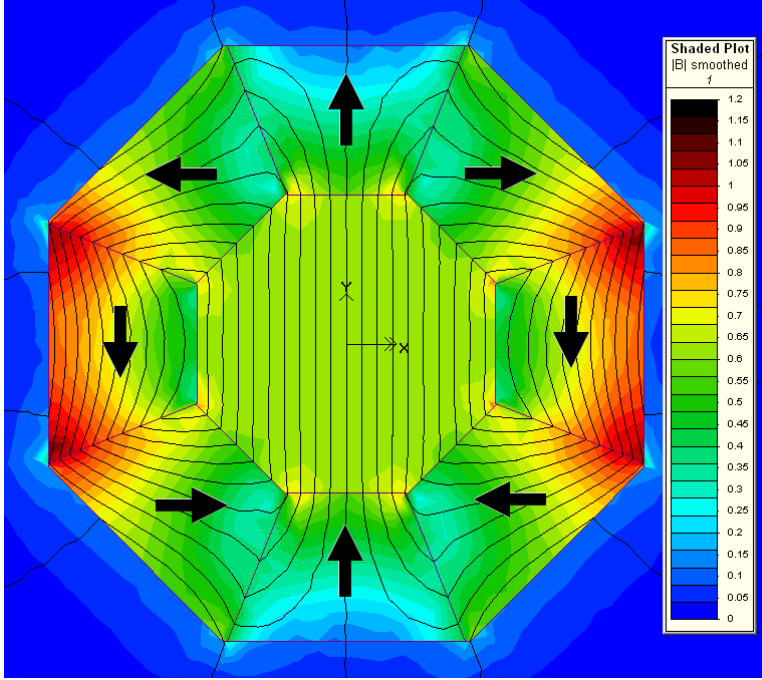
Figure 4-1: $k = 2$ Halbach ring. Lines and arrows show direction of magnetic field, color shows relative magnitude. Image source: Wikipedia

magnetic field noise and to generate a stronger field, an array of rare earth permanent magnets was used. The magnetic field across the diamond needs to be as uniform as possible so that each NV center sees the same field. COMSOL simulations were used to assess the uniformity of various configurations and compare them against the uniformity of the Helmholtz coils. The uniformity was calculated as the deviation of the magnitude of the field divided by its mean. Our smallest Helmholtz coils are 8 inches in diameter, which lead to a simulated uniformity of 0.5% over a 2 cm^3 cube at the center. After simulating a few different configurations, a $k = 2$ Halbach ring was chosen as the best option [13]. This is a hollow ring of magnetic material where the magnetization vector rotates at twice the azimuthal angle along the ring, as shown in Fig. 4-1.

An 8 inch Halbach ring was simulated in COMSOL and resulted in a uniformity of 0.17% over a 2 cm^3 volume at the center (Fig. 4-2a). Another 8 inch ring composed of bar magnets was simulated in COMSOL with a uniformity of 1% (Fig. 4-2b). This is a factor of two worse than the uniformity of the Helmholtz coils, but is much easier
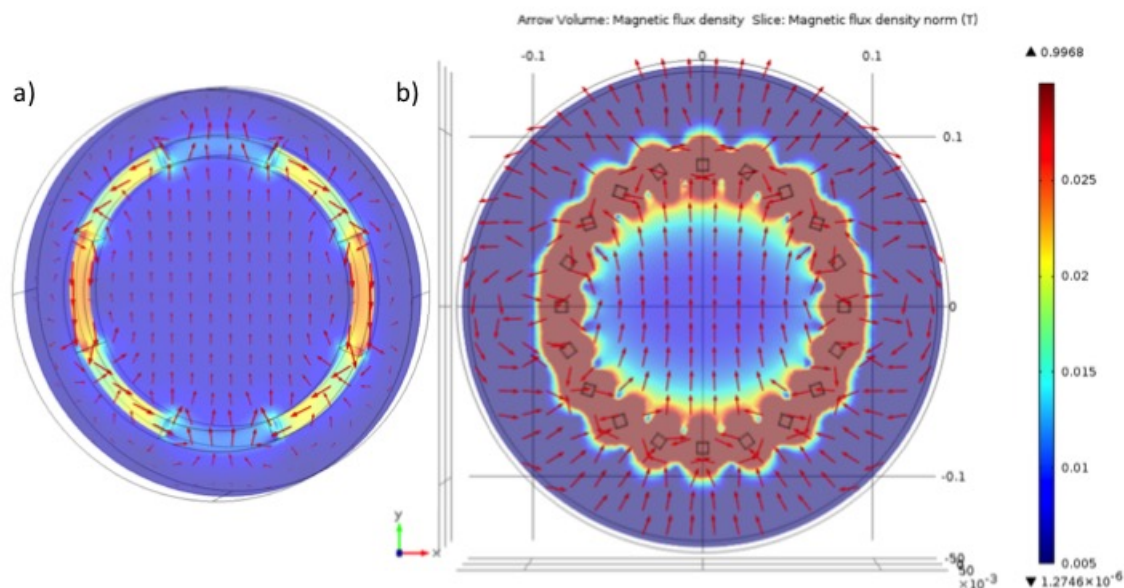
Figure 4-2: a) Circular Halbach ring simulated in COMSOL, arrows show direction
of field, color shows magnitude. b) Halbach ring composed of square bar magnets.

to fabricate than the circular Halbach ring.

This Halbach ring was fabricated by laser cutting an acrylic ring with square holes
at the correct angles for neodymium bar magnets to fit into. The structure produces
a field of 58 G at the center. It is placed inside the Helmholtz coils and can be rotated
to a direction where the eight resonances are equally spaced in frequency. Fitting of
the hyperfine structure of an ESR spectrum resulted in about the same linewidth
(~ 1 MHz) with and without the Halbach ring, indicating a uniformity comparable
to the Helmholtz coils.
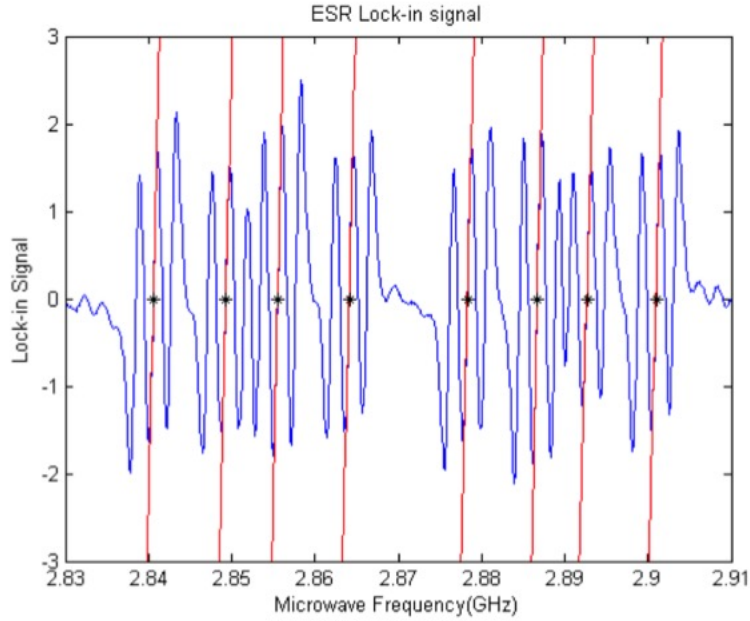
Figure 4-3: Lock-in output of ESR spectrum (blue) with center frequencies marked with black stars and fitted slopes (red lines).

4.2 Sequential Vector Measurements

The direction of the static magnetic field from the Halbach ring and also the orientation of the diamond are determined using the protocol described in the previous chapter. With this information, each resonance can be associated with its crystallographic axis which has a known direction in the lab frame. This is now enough information to measure an external field and reconstruct its magnitude and direction

The protocol for doing this was to first record an ESR spectrum of the lock-in amplifier output and find the center frequencies of each of the eight resonances as well as the slope of the curve nearby the center, as shown in Fig. 4-3

The measurement was then conducted by sequentially hopping between each of the eight recorded frequencies and measuring the lock-in output in the same way as was done for the single axis measurement (see chapter 2), where the slope is used to convert the lock-in voltage to a frequency shift. When all eight frequency shifts have been measured, the vector field is calculated. Fig. 4-4 shows an experiment
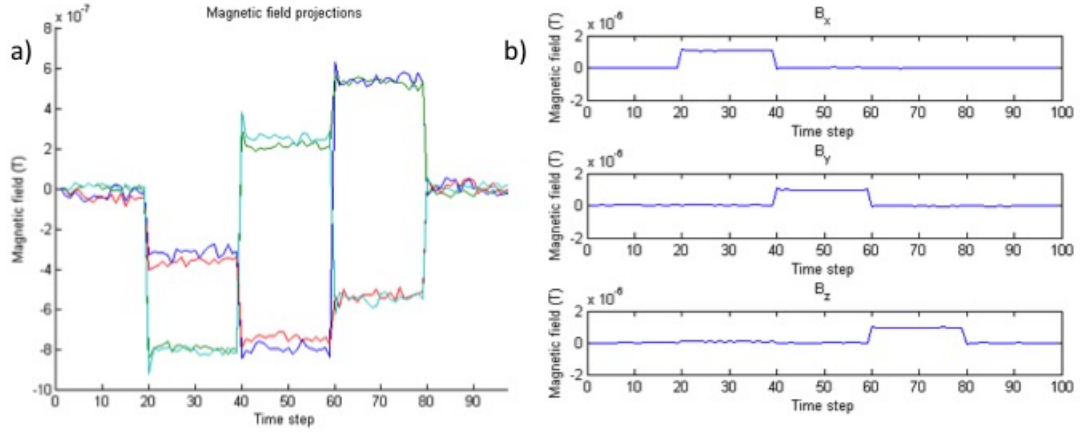
Figure 4-4: A 1 μT step in the magnetic field was applied first in the x direction between time steps 20 and 40, then in the y between 40 and 60, and finally along z between 60 and 80. A time step was the amount of time required to measure all eight frequencies (about .5 seconds). a) The projections of the measured field along each axis b) The reconstructed vector components

where the Helmholtz coils were used to apply 1 μT steps in magnetic field along each direction and the diamond was used to measure this applied field. 4-4a) shows the projections of the field along each axis and 4-4b) shows the reconstructed vector components. Although the diamond was placed at an arbitrary angle with respect to the Helmholtz coils, the direction of the measured vector field is in good agreement. However significant noise (around 50 - 100 nT p-p) limited the sensitivity of this approach.

4.3 Differential measurement

A differential magnetometer or gradiometer measures the magnetic field using two sensors placed some distance apart. The difference in the two measurements gives the gradient of the field between the two points. The advantage of this type of device is the removal of common mode noise seen by both sensors. This same approach was explored with the NV based magnetometer by using two physically separated diamonds. A diagram of the setup is shown in Fig. 4-5. The same 532 nm laser was split into two paths and used to excite both diamonds which were placed about 1 cm apart with a microwave antenna in between. Each diamond was therefore exposed to microwaves of the same frequency and of approximately equal power. The fluorescence of each diamond was collected separately using two detectors.

The two diamonds were not aligned in the same orientation, so a separate rotation matrix had to be found for each. Fig. 4-6 shows the ESR spectra of the two diamonds. Because they are not both in the same orientation, the same static field produces different resonant frequency splittings.

Data is collected in the same sequential manner described in the previous section. Since only one microwave antenna and one source were used, it was necessary to step through all eight resonances of one diamond and subsequently all eight of the other diamond, so that all measurements were done serially. Fig. 4-7a) shows the same stepping experiment as Fig. 4-4 but measured with both diamonds. Despite their different orientations, the vector vector field measured from each is in good agreement. Fig. 4-7b) shows the difference in the components of the two diamonds, or equivalently, the gradient in the direction between the diamonds.

Unfortunately, the noise level in the difference signal is about the same as it is for each measurement individually. This is most likely because higher frequency common mode noise is not being eliminated. For these measurements, the lock-in time constant was set at 10 ms, so that each time the frequency hops to a different resonance, the lock-in takes around 10 ms to settle to its final value. So to cycle through all eight resonances takes about 200 ms and since the diamonds are measured sequentially,
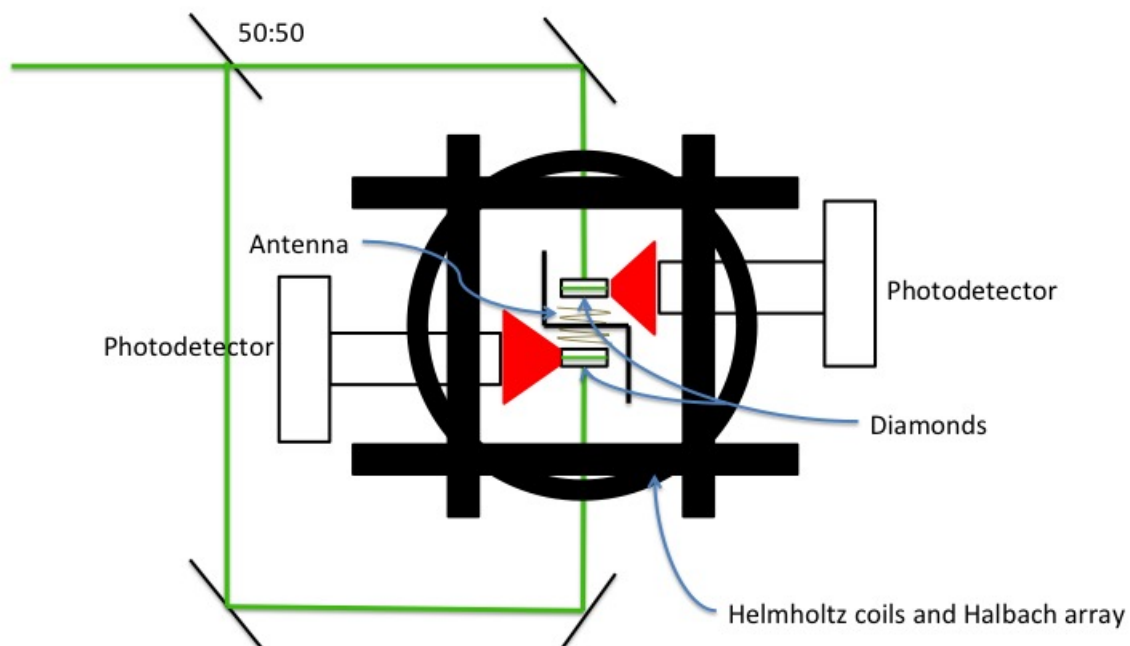
Figure 4-5: Differential Magnetometer setup. The green 532 nm laser is split into to paths of equal power, one illuminating each diamond. The diamonds are placed at the center of the Halbach ring and Helmholtz coils, about 1 cm apart, arranged vertically such that their smallest dimension is in line. A single microwave antenna composed of a wire with three loops, was placed in between the two diamonds. Two photodetectors were used, one for each diamond, with a light block placed in between the diamonds to separate the fluorescence.
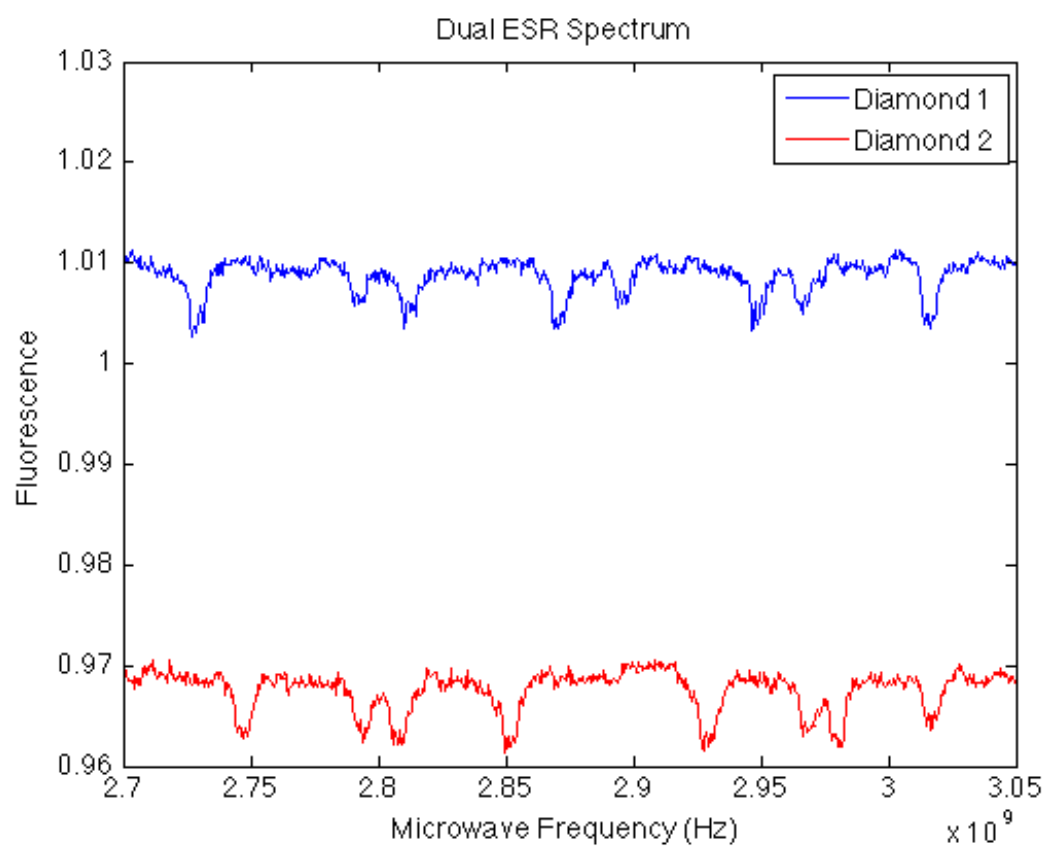
Figure 4-6: ESR spectra of both diamonds
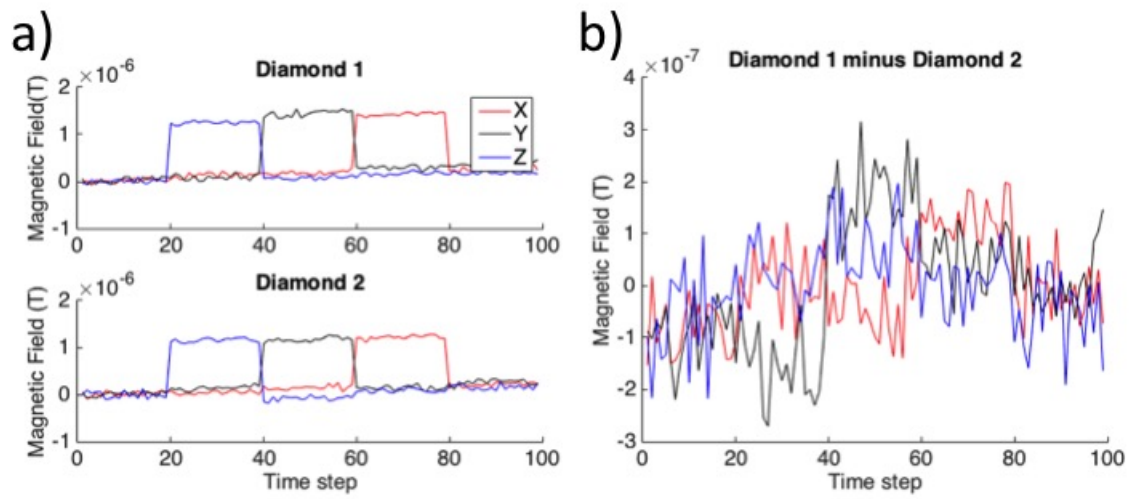
Figure 4-7: a) Vector field components measured by both diamonds for 1 μT steps in x, y, and z. b) Difference between field components of diamond 1 and diamond 2

this is the time between measurements. That means that any noise at frequencies higher than 1/200 ms = 5 Hz will not be common to both measurements and will therefore not be removed. The lock-in time constant could be reduced, but that would allow even more higher frequency noise to come through and the same problem would remain. To properly subtract out the common mode noise, it is necessary to perform the measurements on the two diamonds simultaneously. This could be done for instance by perfectly aligning the two diamonds so that they had the same lattice orientation and their resonant frequencies overlapped.

Instead of taking this approach, a frequency locking scheme was developed that uses dual demodulation to measure shifts in separate resonances at the same time. By employing this scheme on all four orientations, it should be possible to simultaneously measure the vector fields seen by both diamonds.

Chapter 5

Frequency Locking

The frequency locking experiment is a slightly different approach to the measurement of the resonant transition frequency shifts, where a closed-loop feedback system locks a microwave source to each of the ground state transitions, $m_s = 0 \rightarrow 1$ and $m_s = 0 \rightarrow -1$. Then the frequency difference between the two locked sources is recorded, providing a direct measurement of the magnetic field projection on to the NV axis. This technique has so far only been implemented on a single NV orientation, however when extended to all four orientations, it should allow for simultaneous observation of the four projections since a separate microwave source is locked to each transition. Compared with previously implemented, open-loop measurements, this technique also provides far better dynamic range and steady-state immunity to fluctuations in transition linewidth and contrast.

Up to this point, all continuous wave nitrogen vacancy magnetometer measurements (including those described in the previous chapters) have involved an initial calibration measurement which is presumed constant and a second measurement to infer the ambient field using this calibration [2, 1, 15]. For the experiments described in the previous chapters, this calibration was the measurement of the slope of the lock-in output near the resonance frequency. This measurement approach is inherently dependent on other factors besides a true frequency shift, such as the linewidth and contrast of the resonances which may be influenced by factors such as laser pump power, microwave power and detection efficiency [14]. All of these factors are liable

to change from device to device and thus require calibration ahead of time. Some of these parameters might also change over the course of a day, requiring periodic re-calibration. The frequency locking approach provides a direct measurement of the frequency splitting which is tied to fundamental constants instead of a calibration factor

5.1 Implementation

This technique builds upon the work described in previous chapters and in the publication [4]. Using a similar apparatus, we monitor the $m_s = 0 \rightarrow \pm 1$ resonance frequencies of a single NV orientation simultaneously. We use two microwave sources, one locked to the $m_s = 0 \rightarrow \pm 1$ transition, the other to the $m_s = 0 \rightarrow \mp 1$ transition (that is, $m_s = 0 \rightarrow 1$ and $m_s = 0 \rightarrow -1$ for NVs oriented one way and $m_s = 0 \rightarrow -1$ and $m_s = 0 \rightarrow 1$ for NVs oriented along the same crystal axis, but opposite direction). A diagram of the signal path for one of the two frequency locked loops is shown in figure 5-1. The microwave frequency generator is used as a tunable source with its frequency modulation input connected to the output of a feedback compensator.

For purpose of dual demodulation, each source is frequency modulated at a different rate with a modulation depth of 260 kHz. Using two different modulation frequencies allows us to observe changes in fluorescence at the two different microwave frequencies on the ESR curve. This is done using an IQ mixer which adds a binary frequency shift modulation at a rate (f_{ref}) of 3 kHz for one loop and 4 kHz for the other, each with a modulation depth (f_{dev}) of 260 kHz. The modulation depth is chosen to be a little less than the width of the hyperfine resonance as this gives the largest signal. The output of the IQ mixer is then a signal whose frequency jumps between f_{LO}(around 2.8 GHz) $- f_{dev}$ and $f_{LO} + f_{dev}$ at a rate of f_{ref}.

This signal is sent to an antenna which radiates the microwaves over the diamond sample. The fluorescence from the diamond is detected with a photodetector and dual demodulated by a pair of lock-in amplifiers whose internal reference oscillators provide the modulation frequencies (f_{ref}) for the microwave sources. The frequency
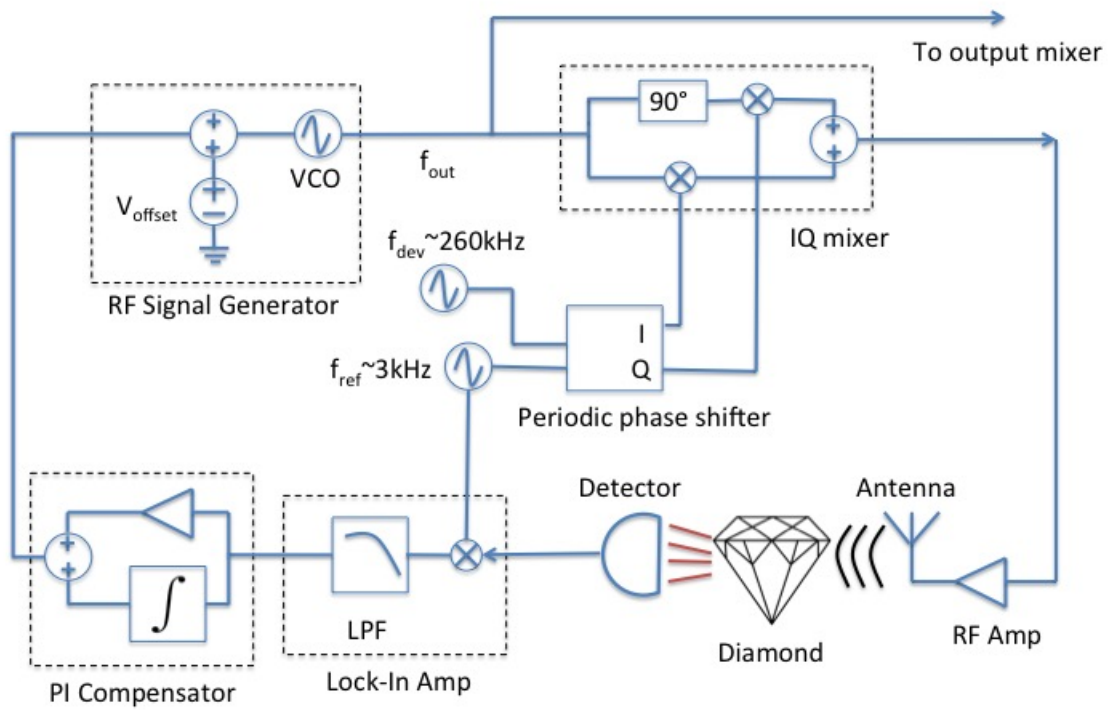
Figure 5-1: Signal path for one of the two frequency locked loops

modulation of the microwaves is converted into an amplitude modulation of the fluorescence due to the frequency dependent shape of the ESR curve. A larger slope in the ESR curve generates a larger amplitude modulation. The amplitude modulated signal is demodulated by the lock-in amplifier which outputs the envelope of the modulation which is proportional to the derivative of the original ESR curve. Thus the lock-in amplifiers serve multiple purposes: firstly, they eliminate detector noise, secondly, they allow us to simultaneously observe frequency shifts of two different transitions, and finally, they provide us with a derivative signal which is zero at the center of the resonance and linear on either side.

The output of the lock-in amplifiers provides an error signal which is zero if the microwave source is on resonance, positive if the source is too high, and negative if the source is too low. This error signal is fed into a PI compensator whose output goes to the frequency modulation port of the microwave source, closing the feedback loop. The difference frequency is readout by mixing the output of the two microwave frequency generators. The mixed signal is measured with a frequency counter with a bandwidth lower than the sum frequency of the two sources, such that only the difference frequency is measured.

5.2 Experimental Results

Adjusting the gain or cutoff frequency of the PI compensator changes the system's dynamic response while maintaining the same steady-state value. In particular, the integrator term allows us to boost the low frequency gain and eliminate steady-state error while preserving stability. Similarly, the steady-state response of the system is immune to changes in the slope of the lock-in signal, which is equivalent to changes in the gain of the compensator. These changes in slope come about due to variations in the resonance linewidth and contrast. In Figure 5-2a) the slope of the lock-in signal is changed by varying the laser pump power which changes the contrast of the ESR curve. Figure 5-2b) demonstrates how different pump powers have different transient responses to a step input, but each case converges to the same steady-state result.
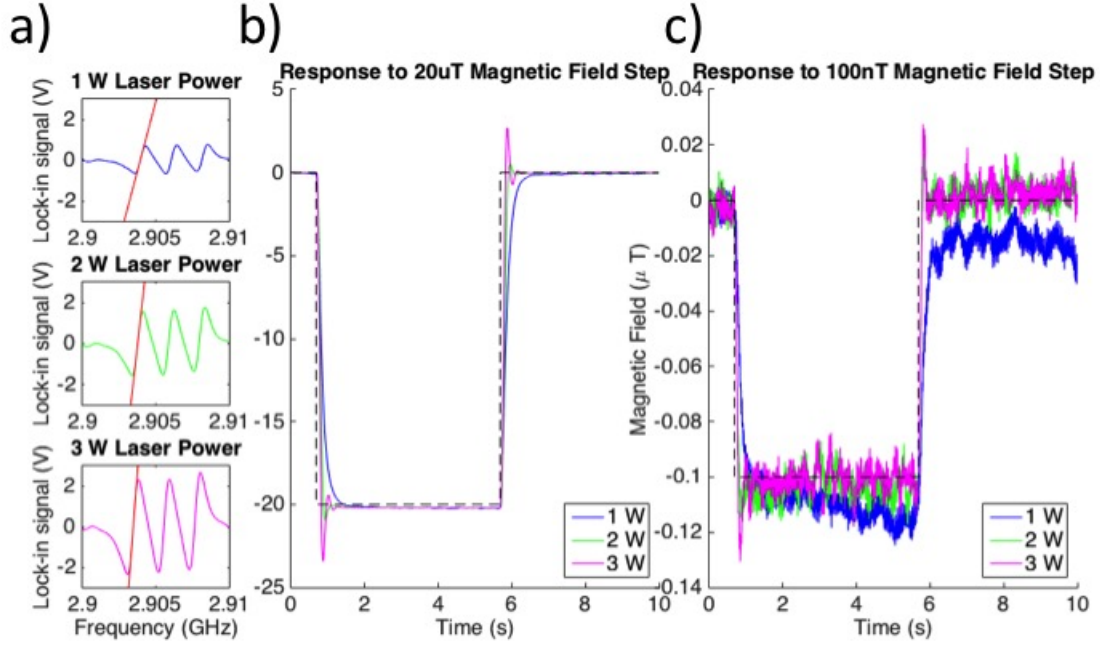
Figure 5-2: Step response dependence on laser power. a) The pump laser power was increased from 1W to 2W to 3W resulting in a steeper lock-in slope which is equivalent to a larger gain constant of the PI compensator. b) The response to a 20 uT step in magnetic field for each pump power is shown. The dynamics of the response are changed but the steady-state values are identical. The settling time for each is around 200 ms. c) The response to a 100 nT step in magnetic field for the same three laser pump powers.

Figure 5-2b) also demonstrates the increased dynamic range of this approach where a step magnetic field of 20 μT is applied. This corresponds to a 560 kHz shift which extends beyond the limits of the linear region of the lock-in signal. The dynamic range is only limited in our system by the 4 MHz tuning range of our microwave sources.

Figure 5-3 shows the noise power spectral density and corresponding Allan deviation plot with a 5 nT, 1 Hz applied calibration field. The sensitivity of the device is given by the amplitude of the noise at each frequency. For frequencies below around 2 Hz, the sensitivity is around 1 nT Hz$^{-1/2}$, which is comparable to our previous, open
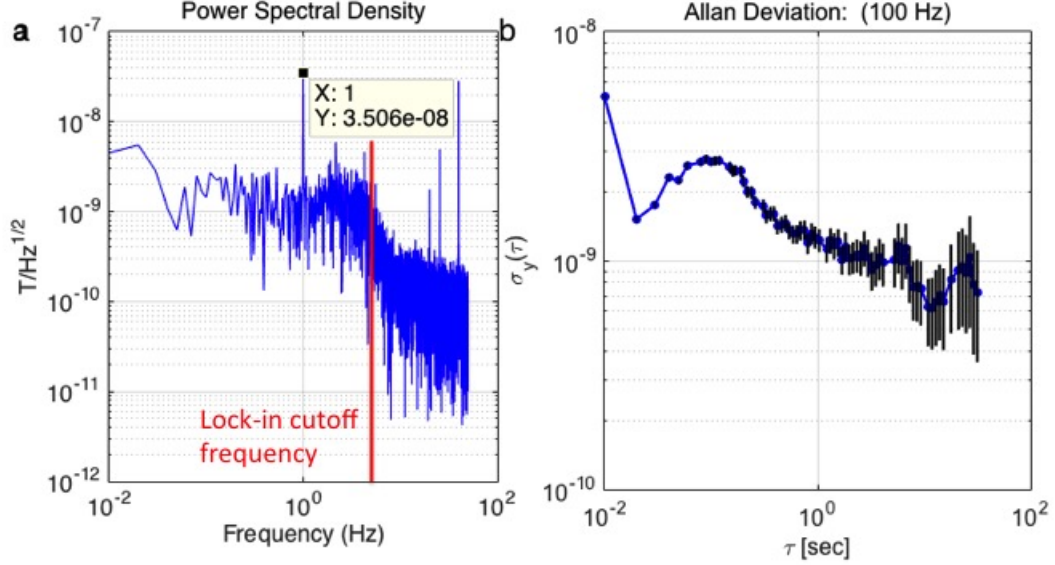
Figure 5-3: a) Noise Power Spectral Density with the lock-in cutoff frequency at the red line and b) corresponding Allan Deviation plot with a 5 nT reference field applied at 1 Hz

loop implementation. The power spectral density was calculated by sampling the measured magnetic field at a rate of 100 Hz for 100 s. This results in a frequency resolution of .01 Hz and the 5 nT reference field generates the $5 \text{ nT}/(\sqrt{2}\sqrt{.01 \text{ Hz}}) = 35 \text{ nT/Hz}^{1/2}$ peak at 1 Hz. The Allan deviation was calculated using the same time trace. It shows a continual decrease in the deviation between subsequent measurements as the integration time is increased, indicating that we have not yet run up against thermal or flicker noise for these time scales.

5.3   Details of frequency modulation circuit

In order to accurately count the frequency difference between the two sources, it was necessary to apply the frequency modulation external to the source, allowing the
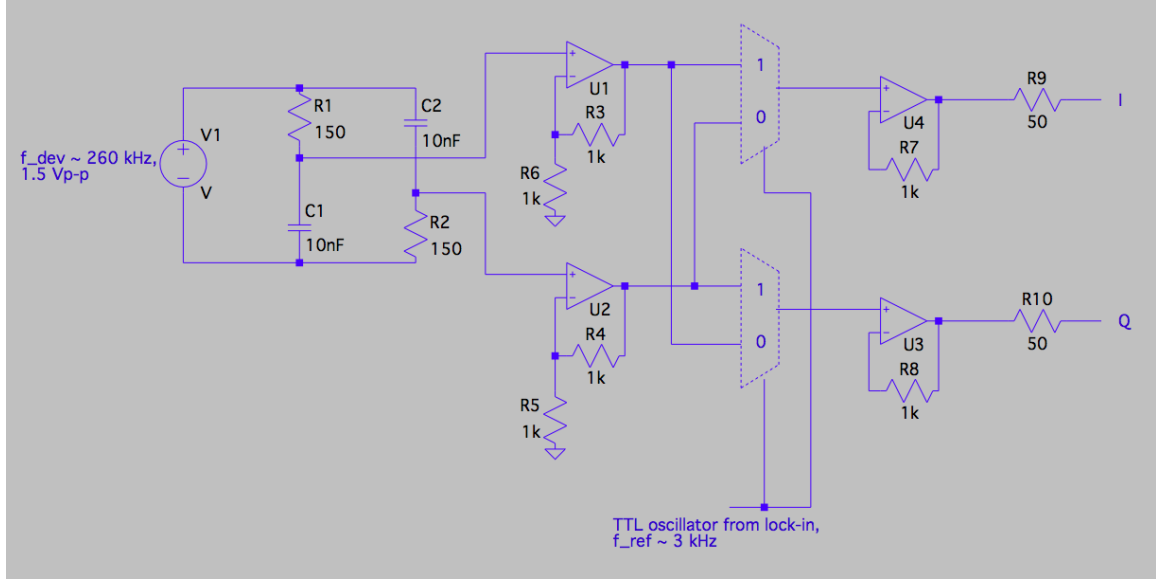
Figure 5-4: Schematic of phase shifter circuit to drive I and Q inputs of IQ mixer

pure tone to be measured first. This is done using an IQ mixer, as described above, where the I and Q ports are connected to a periodic phase shifter that outputs two sinusoidal signals at the deviation frequency ($f_{dev} \approx 260$ kHz) with a phase difference that oscillates between $+90^o$ and -90^o at a rate of $f_{ref} \approx 3$ kHz. The phase shifter circuit is shown in fig. 5-4.

The two RC dividers formed by R1,C1 and R2,C2 output signals that are 90^o out of phase. The values are chosen such that the magnitude of the impedance of the capacitance and resistance are the same at 260 kHz ($1/(2\pi RC) \approx f_{dev} \approx 260$ kHz). Op amps U1 and U2 buffer the signals and amplify by a factor of 2. Two multiplexers are used to control the connections between the two phase shifted signals and the I and Q outputs. U3 and U4 serve to buffer the signal (since the multiplexers have high output impedance) and drive the 50 Ohm load of the IQ mixer. The multiplexers' inputs are connected to the TTL oscillator from the lock-in amplifier. Fig. 5-5 shows the I and Q waveforms captured using an oscilloscope.
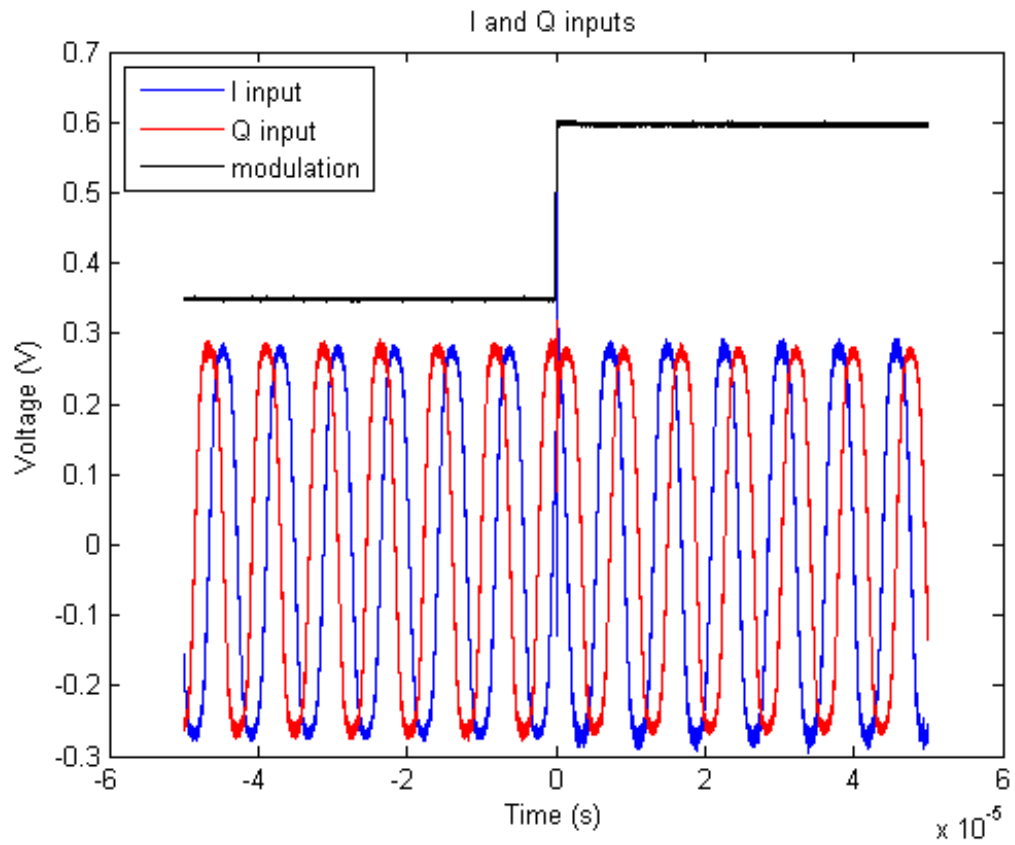
51

Figure 5-5: I (blue) and Q (red) waveforms. Black line labeled modulation is the reference oscillator from the lock-in. When this signal is low, I leads Q by 90^o and when it's high, I lags Q by 90^o.

Chapter 6

Conclusion

Our experiments thus far have demonstrated competitive sensitivities of around $1\,\mathrm{nT}/\sqrt{\mathrm{Hz}}$ for measurements which are capable of separating temperature and magnetic field effects from resonant frequency shifts. The development of a procedure to determine the diamond lattice orientation has enabled preliminary vector measurements with good agreement between measured and applied fields. A differential magnetometer was investigated in an attempt to reduce measurement noise but was limited to cancelling only very low frequencies due to the sequential measurement technique. The frequency locking technique improves performance by being immune to variations in resonance line-shape, greatly increasing the achievable dynamic range, and allowing for simultaneous measurements of more than one axis.

6.1   Future Experiments

Logically, the next challenge would be to implement the frequency locking technique on all four NV orientations. This is most likely a technical challenge more than anything else. The development of some custom hardware would be necessary to reduce the number of required instruments, which is already quite substantial for just a single axis. For all four orientations, eight separate lock-in amplifiers, feedback compensators, and microwave sources and modulators would be required. The lock-in amplifiers and feedback compensators could be implemented digitally on a single

FPGA while the microwave sources and modulators can be found as inexpensive integrated circuits that could be implemented on a single PCB.

Once the frequency locking, vector magnetometer was demonstrated, another differential experiment could be done where vector measurements would be performed simultaneously on two separate diamonds and then subtracted to find the gradient, suppressing all common mode noise. This experiment has the potential to significantly reduce the device's noise floor, removing common mode magnetic and instrument noise.

The frequency locking technique could also be implemented as a pulsed experiment. A Ramsey type sequence could be applied simultaneously to each of the two transitions. The free precession time would be fixed, and the detuning of the microwave source from resonance would be held constant by locking to a maximum or minimum of the Ramsey fringe (where the derivative would be zero). The two sources would be frequency modulated at a rate much slower than the Ramsey sequence, and the fluorescence signal demodulated using two lock-in amplifiers. This pulsed measurement should result in even better sensitivity due to the increased contrast and potentially narrower linewidth.

Additional improvements in the sensitivity of NV vector magnetometers through such experiments in combination with device miniaturization will open up potential applications in the areas of object detection, navigation and biomedical sensing.

Bibliography

[1] V. M. Acosta, E. Bauch, M. P. Ledbetter, a. Waxman, L. S. Bouchard, and D. Budker. Temperature dependence of the nitrogen-vacancy magnetic resonance in diamond. *Physical Review Letters*, 104(7):1–4, 2010.

[2] Gopalakrishnan Balasubramanian, I Y Chan, Roman Kolesov, Mohannad Al-Hmoud, Julia Tisler, Chang Shin, Changdong Kim, Aleksander Wojcik, Philip R Hemmer, Anke Krueger, Tobias Hanke, Alfred Leitenstorfer, Rudolf Bratschitsch, Fedor Jelezko, and Jörg Wrachtrup. Nanoscale imaging magnetometry with diamond spins under ambient conditions. *Nature*, 455(7213):648–651, 2008.

[3] P Brauer, T Risbo, J.M.G Merayo, and O.V Nielsen. Fluxgate sensor for the vector magnetometer onboard the 'Astrid-2' satellite. *Sensors and Actuators A: Physical*, 81(1-3):184–188, 2000.

[4] Hannah Clevenson, Matthew E. Trusheim, Carson Teale, Tim Schröder, Danielle Braje, and Dirk Englund. Broadband magnetometry and temperature sensing with a light-trapping diamond waveguide. *Nature Physics*, 11(May):393–397, 2015.

[5] Marcus W. Doherty, Neil B. Manson, Paul Delaney, Fedor Jelezko, JÃűrg Wrachtrup, and Lloyd C.L. Hollenberg. The nitrogen-vacancy colour centre in diamond. *Physics Reports*, 528(1):1 – 45, 2013. The nitrogen-vacancy colour centre in diamond.

[6] Florian Dolde, Helmut Fedder, Marcus W. Doherty, Tobias Nöbauer, Florian Rempp, Gopalakrishnan Balasubramanian, Thomas Wolf, Friedemann Reinhard, Lloyd C. L. Hollenberg, Fedor Jelezko, and Jörg Wrachtrup. Sensing electric fields using single diamond spins. *Nature Physics*, 7(6):459–463, 2011.

[7] a. Dréau, M. Lesik, L. Rondin, P. Spinicelli, O. Arcizet, J. F. Roch, and V. Jacques. Avoiding power broadening in optically detected magnetic resonance of single NV defects for enhanced dc magnetic field sensitivity. *Physical Review B - Condensed Matter and Materials Physics*, 84(19):1–8, 2011.

[8] Alan Edelstein. Advances in magnetometry. *Journal of Physics: Condensed Matter*, 19(16):165217, 2007.

[9] a. M. Edmonds, U. F S D'Haenens-Johansson, R. J. Cruddace, M. E. Newton, K. M C Fu, C. Santori, R. G. Beausoleil, D. J. Twitchen, and M. L. Markham. Production of oriented nitrogen-vacancy color centers in synthetic diamond. *Physical Review B - Condensed Matter and Materials Physics*, 86(3):1–7, 2012.

[10] Kejie Fang, Victor Acosta, Charles Santori, Zhihong Huang, Kohei Itoh, Hideyuki Watanabe, Shinichi Shikata, and Raymond Beausoleil. High-Sensitivity Magnetometry Based on Quantum Beats in Diamond Nitrogen-Vacancy Centers. *Physical Review Letters*, 110(13):130802, March 2013.

[11] S. Felton, a. Edmonds, M. Newton, P. Martineau, D. Fisher, D. Twitchen, and J. Baker. Hyperfine interaction in the ground state of the negatively charged nitrogen vacancy center in diamond. *Physical Review B*, 79(7):075203, February 2009.

[12] K. M C Fu, C. Santori, P. E. Barclay, and R. G. Beausoleil. Conversion of neutral nitrogen-vacancy centers to negatively charged nitrogen-vacancy centers through selective oxidation. *Applied Physics Letters*, 96(12):2013–2016, 2010.

[13] K. Halbach. Design of permanent multipole magnets with oriented rare earth cobalt material. *Nuclear Instruments and Methods*, 169(1):1–10, 1980.

[14] K. Jensen, V. M. Acosta, a. Jarmola, and D. Budker. Light narrowing of magnetic resonances in ensembles of nitrogen-vacancy centers in diamond. *Physical Review B - Condensed Matter and Materials Physics*, 87(1):1–10, 2013.

[15] K Jensen, N Leefer, a Jarmola, Y Dumeige, V M Acosta, P Kehayias, B Patton, and D Budker. Cavity-enhanced room-temperature magnetometry using absorption by nitrogen-vacancy centers in diamond. *ArXiv*, 160802(April):5, 2014.

[16] I K Kominis, T W Kornack, J C Allred, and M V Romalis. A subfemtotesla multichannel atomic magnetometer. *Nature*, 422(6932):596–599, 2003.

[17] D. Le Sage, L. Pham, N. Bar-Gill, C. Belthangady, M. Lukin, a. Yacoby, and R. Walsworth. Efficient photon detection from color centers in a diamond optical waveguide. *Physical Review B*, 85(12):1–4, 2012.

[18] J. Lenz and S. Edelstein. Magnetic sensors and their applications. *IEEE Sensors Journal*, 6(3):631–649, 2006.

[19] James E. Lenz. Review of magnetic sensors. *Proceedings of the IEEE*, 78(6):973–989, 1990.

[20] J R Maze, P L Stanwix, J S Hodges, S Hong, J M Taylor, P Cappellaro, L Jiang, M V Gurudev Dutt, E Togan, a S Zibrov, a Yacoby, R L Walsworth, and M D Lukin. Nanoscale magnetic sensing with an individual electronic spin in diamond. *Nature*, 455(7213):644–647, 2008.

[21] Pavel Ripka. Advances in fluxgate sensors. *Sensors and Actuators, A: Physical*, 106(1-3):8–14, 2003.

[22] Lucio Robledo, Hannes Bernien, Toeno Van Der Sar, and Ronald Hanson. Spin dynamics in the optical cycle of single nitrogen-vacancy centres in diamond. *New Journal of Physics*, 13, 2011.

[23] Yumin Shen, Timothy M. Sweeney, and Hailin Wang. Zero-phonon linewidth of single nitrogen vacancy centers in diamond nanocrystals. *Physical Review B - Condensed Matter and Materials Physics*, 77(3):2007–2009, 2008.

[24] J. M. Taylor, P. Cappellaro, L. Childress, L. Jiang, D. Budker, P. R. Hemmer, a. Yacoby, R. Walsworth, and M. D. Lukin. High-sensitivity diamond magnetometer with nanoscale resolution. *Nature Physics*, 4(10):810–816, September 2008.